Transformed Residual Quantization for Approximate Nearest Neighbor Search

Jiangbo Yuan
Florida State University
jyuan@cs.fsu.edu

Xiuwen Liu
Florida State University
liux@cs.fsu.edu

arXiv:1512.06925v1 [cs.CV] 22 Dec 2015

Abstract

The success of product quantization (PQ) for fast nearest neighbor search depends on the exponentially reduced complexities of both storage and computation with respect to the codebook size. Recent efforts have been focused on employing sophisticated optimization strategies, or seeking more effective models. Residual quantization (RQ) is such an alternative that holds the same property as PQ in terms of the aforementioned complexities. In addition to being a direct replacement of PQ, hybrids of PQ and RQ can yield more gains for approximate nearest neighbor search. This motivated us to propose a novel approach to optimizing RQ and the related hybrid models. With an observation of the general randomness increase in a residual space, we propose a new strategy that jointly learns a local transformation per residual cluster with an ultimate goal to reduce overall quantization errors. We have shown that our approach can achieve significantly better accuracy on nearest neighbor search than both the original and the optimized PQ on several very large scale benchmarks.

1. Introduction

Nearest neighbor search in very large databases is becoming increasingly important in machine learning, computer vision, pattern recognition, and multimedia retrieval along with many applications in document, image, audio, and video retrievals [23, 33, 25, 9, 1, 27, 30, 5, 18, 17, 15]. However, it becomes difficult to efficiently store and search huge collections when the dataset size gets larger and larger (e.g. millions or even billions). The idea of mapping real-valued vectors to compact codes [31, 19, 14] provides a very attactive way to solve this problem. There have been many recent methods on developing hashing methods for compressing data to compact binary strings [31, 14, 1, 22, 7, 32, 24], or vector quantization based methods [19, 11, 28] that compress data to compact codes.

Vector quantization (VQ) was actively studied for source coding/signal compression under real-time constraints dating back to decades ago [13]. Recently, the problem of how to apply VQ techniques to efficient approximate nearest neighbor (ANN) search has attracted a lot of attension [19, 8, 35, 3, 11, 28, 37, 2]. While many structured VQ models can be found [13, 16, 4], in this paper, we restrict our attention to product quantization (PQ) and residual quantization (RQ), which both have been successfully applied to fast nearest neighbor search [19, 8].

1.1. Related Works

Product quantization works by grouping the feature dimensions into groups, and performs quantization to each feature group. In particular, it performs a k-means clustering to each group to obtain sub-codebooks, and the global quantization codebook is generated by the Cartesian products of all the small sub-codebooks. In this way, it can generate a huge number of landmark points in the space, which guarantees low quantization error; it has achieved state of the art performance on approximate nearest neighbor search [19], and can also provide a compact representation to the vectors. Inspired by the success of PQ, some latest works have extended PQ to a more general model by finding an optimized space-decomposition to minimize its overall distortion [11, 28]. A very recent work [21] has deployed this optimized PQ within residual clusters. While it maximizes the strength of locality, it also uses extra space for multiple transformations as well as PQ codebooks.

Different from PQ, RQ works by performing quantization to the whole feature space, and then recursively apply VQ models to the residuals of the previous quantization level, which is a stacked quantization model. In particular, it performs a k-means clustering to the original feature vectors, and construct k clusters. For points in each cluster, it computes the residuals between points and the cluster centers. In the next level, it aggregates all the residual vectors for all points, and performs another clustering to these residual vectors. This process is recursively applied (stacked) for several levels. In this way, RQ produces sequential-product codebooks. A comprehensive survey of earlier RQ models can be found in Barnes et al. [4]. Recent works have shown the effectiveness of RQ to perform both indexing [35] and data compression [8] tasks in ANN search problems.

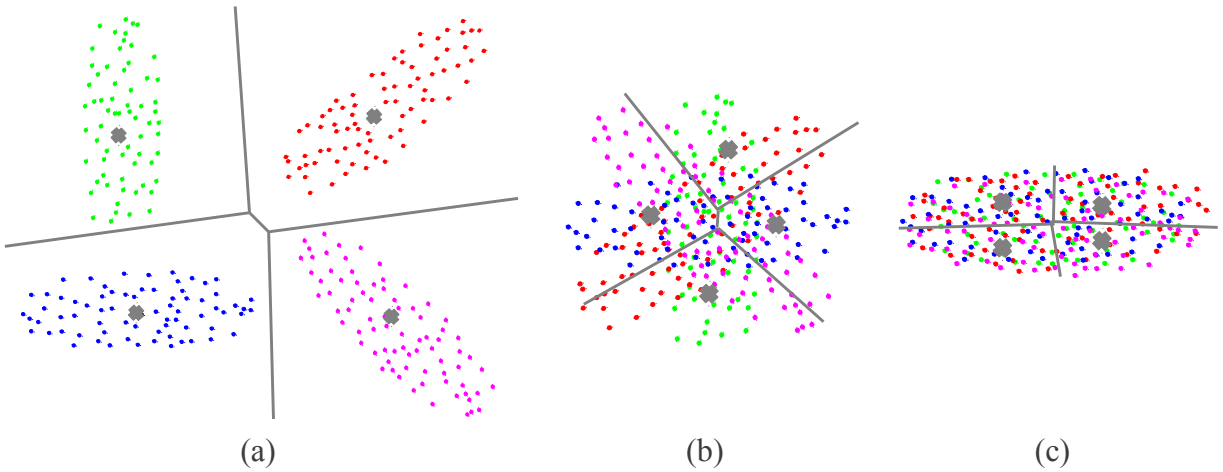

Figure 1: A toy example of two-level RQ and TRQ. While the first-level vector quantizer (a) is identical to the two models, the residual space in the ordinary RQ model (b) is much noisy than the TRQ model (c), in which the residual clusters are aligned by our novel method.

However, it has been observed that the effectiveness of RQ might be limited to a very small number of levels [13], and the randomness of residual vectors increases very quickly when we stack more and more RQ layers. This results in increasingly noisy residual vectors of each level, and effective quantization at higher RQ levels becomes more difficult. In other words, the layer-wise learning in RQ models provides suboptimal sub-codebooks in that each sub-codebook is learned without consideration to subsequent layers. Thence several global optimization approaches have been proposed by jointly learning sub-codebooks over all layers [6, 37, 2]. Different from seeking out joint optimization solutions, we here look for specific transformations to each of the residual clusters to make the resulting vectors more aligned, in order to directly address the problem of cluster misalginment and noise increase over each level after the first and improve overall quantization accordingly.

1.2. Contributions

To this end, we propose a novel approach to optimizing the RQ model that is motivated by the noise and shape in the residual space, as shown in Figure 1. We find that the residual vectors for each cluster have significantly different directions, which make quantization at the next level much harder. A natural idea is to align these residual vectors for each cluster, which can potentially reduce quantization error in the next level. Thus, we propose to learn one rotation matrix for each residual cluster, and use them to align the residual vectors in each cluster, so as to reduce the global quantization error. Then we alternate between learning rotation matrices and the residual quantization to minimize its distortion, which is mainly inspired by "Iterative Quantization (ITQ)" [14]; very recently ITQ has been successfully applied to the PQ models [11, 28, 21]. Different from ITQ and the optimized PQ (OPQ), which learn a global rotation, both ours and the locally optimized PQ (LOPQ) [21] learn one projection matrix per residual cluster. In stead of independent learning in LOPQ, however, all local transformations in our method are learned associatively with a shared codebook. In this way, we successfully reduce the memory overheads about codebook usage in LOPQ. This offers a great flexibility to use reasonable larger codebooks in potential.

We have found that by doing this iterative alignment and quantization, we can achieve significantly smaller quantization error than both vanilla RQ and OPQ methods. In addition, we also propose a hybrid ANN search method which is based on the proposed TRQ and PQ. Experimental results on several large-scale datasets have clearly demonstrated the effectiveness of our proposed methods, in particular have shown that the extra transformations only introduced very limited computation overheads when integrated with advanced indexing structures, e.g. the inverted mutli-index [3]. Therefore, our method is able to achieve very significant gain over other state of the art methods in terms of trade-off among memory usage, search quality and speed.

2. Background and Formulations

In this section, we review related background on two types of structured vector quantization, i.e. product quantization and residual quantization. To accurately quantize large number of points in a high dimensional space, we need a large number of landmark points (or centroids). For ex-

ample, if we use the k-means clustering method to find such quantizers, when the number of cluster centers k becomes huge (e.g. millions or even more), performing k-means becomes prohibitively expensive. Structured vector quantization make special assumptions about the data distribution, and try to explore such structure to generate large number of landmark points to increase the quantization accurateness [13]. PQ and RQ are two structured VQ families with different assumptions about data distribution. We here present discussions on structural codebook constructions and objective function formulations of the related VQ models following a brief introduction of the unstructured VQ.

Vector quantization (VQ a.k.a. k-means) without any structure constraints is probably one of the most widely used vector quantization method. Given a dataset $X = \{\boldsymbol{x}_j : \boldsymbol{x}_j \in \mathbb{R}^D, j = 1, ..., N\}$, VQ is a mapping:

$$q(\boldsymbol{x}_j) = \boldsymbol{c}_i, \boldsymbol{c}_i \in C \tag{1}$$

where q is a quantizer and $\boldsymbol{c}_i, i = 1, 2, ..., k$ is a *centroid* or a *center* from the codebook C. According to Lloyd's optimality conditions, an optimal quantizer satisfies the nearest neighbor condition: $q(\boldsymbol{x}_j) = \arg\min_{\boldsymbol{c}_i \in C} d(\boldsymbol{x}_j, \boldsymbol{c}_i)$. Here $d(.,.)$ is the distance between two vectors and Euclidean distance is used in this paper. For each centroid, a set of data points will be assigned to it, and forms a cluster. According to the second optimal condition, a centroid is computed as: $\boldsymbol{c}_i = E(\boldsymbol{x}_j | \boldsymbol{x}_j \in V_i)$. We estimate the codebook centers to minimize the objective function: the mean squared error (MSE)

$$\text{MSE}(q) = 1/N \sum_{j=1}^{N} \|\boldsymbol{x}_j - q(\boldsymbol{x}_j)\|_2^2 . \tag{2}$$

While the globally optimal solution of the above problem is NP-hard, it can be solved by heuristic alternatives. The best known approach is the k-means algorithm [13], in which the above two conditions are optimized alternatively.

2.1. Product Quantization

Product quantization assumes that certain groups of features are independent with each other, and explores this assumption to generate a large number of landmark points by grouping the feature space into m groups (each group is a subspace). By quantizing each of the subspaces separately using sub-quantizer $q_i(\cdot)$, it produces an implicit codebook as a Cartesian product of m small sub-codebooks, $C = C_1 \times C_2 \ldots \times C_m$. In this case, codebook C can provide an exponentially large number of cluster centers while retains a linear size of storage. Given a data point $\boldsymbol{x}_j$, it estimates the globally nearest center by concatenating all its nearest sub-centers from the sub-codebooks as $q(\boldsymbol{x}_j) = \boldsymbol{c}_{i1} | \ldots | \boldsymbol{c}_{im}$. The MSE of this product quantizer can be estimated as:

$$\text{MSE}(q) = \sum_{i=1}^{m} \text{MSE}(q_i). \tag{3}$$

The degradation of PQ performance can be severe if there are substantial statistical dependences among the feature groups [13]; Recent works have shown that such dependences could be reduced, to some extent, by more careful space-decomposition [11, 28]. Their works extended the idea from iterative quantization [14] to the PQ scheme that learn an rotation to transform the data to reduce the dependences between feature groups. For example, the work from [11] jointly seeks a whole space rotation T and PQ codebooks C by minimizing

$$\begin{aligned} \text{MSE}(q) &= 1/N \sum_{j=1}^{N} \|T\boldsymbol{x}_j - q(T\boldsymbol{x}_j)\|_2^2 , \\ s.t. \quad & q(T\boldsymbol{x}_j) = T(\boldsymbol{c}_{i_1} | \boldsymbol{c}_{i_2} | ... | \boldsymbol{c}_{i_m}). \end{aligned} \tag{4}$$

2.2. Residual Quantization

Residual quantization (RQ) has a different assumption in that it does not assume the features are independent, but assumes the quantization residuals of the first level quantizer can be further quantized. Thus, it is a stacked quantization model. For the first level, RQ simply uses a k-means clustering to quantize the data, and assign them to k centers. For each data point, by subtracting from the assigned centroid, we can collect their residual vectors as $\boldsymbol{r}_j = \boldsymbol{x}_j - q_1(\boldsymbol{x}_j)$. Then $R = \{\boldsymbol{r}_j, j = 1, ..., N\}$ will be used as the input to the next level, and we again use k-means[1] to quantize the residual vectors. By repeating this for h times, we can have a sequential product codebook $C = C_1 \times C_2 \ldots \times C_h$. Given an input vector $\boldsymbol{x}_j$, by computing the nearest center at each level, we can get a sequence of indexes $(i_1, i_2, \ldots, i_h)$. The globally nearest center for $\boldsymbol{x}_j$ here becomes a direct sum of all the sub-centers $q(\boldsymbol{x}_i) = \sum_{l=1}^{h} \boldsymbol{c}_l$.

Simply considering one level, we represent $\boldsymbol{x}_j$ at level l as $\boldsymbol{x}_j^l$ where $\boldsymbol{x}_j = \boldsymbol{x}_j^1$. Then the residual vector from level l is

$$\boldsymbol{x}_j^{l+1} = \boldsymbol{x}_j^l - q_l(\boldsymbol{x}_j^l), \tag{5}$$

the MSE for a h-level RQ is

$$\begin{aligned} \text{MSE}(q) = \text{MSE}(q_h) &= 1/N \sum_{j=1}^{N} \|\boldsymbol{x}_j^h - q_h(\boldsymbol{x}_j^h)\|_2^2 \\ &= 1/N \sum_{j=1}^{N} \|\boldsymbol{x}_j^{h+1}\|_2^2. \end{aligned} \tag{6}$$

[1] Although other quantizers can be used, such as a product quantizer, using k-means is a basic choice.

The most important advantage of RQ is that it does not make the unrealistic assumption that the features are statistically independent. In addition, it holds a non-increasing property, that is, adding a level will always reduce the MSE error.

3. Transformed Residual Quantization

As discussed above, residual quantization recursively performs quantization on residual vectors from previous levels. In other words, all the residual vectors from different clusters are collected and fed into the same quantizer in the next level. As shown in Figure 1 (b), the residual clusters can have different shapes, orientations, or scales, which makes effective quantization of them very hard in general. To address this problem, we propose to learn one rotation per cluster to transform the residual clusters, to make them more aligned to each other, and potentially can make more fruitful residual-level quantizers (see Figure 1(c)).

3.1. Residual Quantization with Cluster-Wise Transforms

For simplicity, we focus our analysis on a two-level RQ model. We fix the first level to be an k-means clustering with k_1 clusters. And at this point, we are not going to specify the second-level quantizer q_2 considering our proposed method is general to any quantizer performed in the residual space. In this way, we have a first-level codebook C^1 with k_1 cluster centers $\boldsymbol{c}_i^1, i = 1, ..., k_1$. We suppose that a vector $\boldsymbol{x}_j$ is assigned to $\boldsymbol{c}_i^1$, its residual vector is denoted as

$$\boldsymbol{r}_j = \boldsymbol{x}_j - \boldsymbol{c}_i^1, \tag{7}$$

A residual cluster V_i then is a collection of all residual vectors assigned to its center,

$$V_i = \{\boldsymbol{r}_j : q_1(\boldsymbol{x}_j) = \boldsymbol{c}_i^1\}. \tag{8}$$

The residual set $R = \{\boldsymbol{r}_j : j = 1, ..., N\}$ forms the input for the next level quantizer q_2. Finally, the objective function for such a RQ model simply follows either Equation 6 or Equation 3, depending on whether another k-means or PQ is selected to be the residual quantizer. To alleviate the misalignment of these clusters, we propose to apply cluster-specified transforms right after the first level so that all clusters are aligned for better adaption to the shared residual quantizer. We associate each cluster center $\boldsymbol{c}_i^1$ with a transform T_i, whose inverse $(T_i)^{-1}$ must exist. Accordingly, a transformed residual cluster is then represented as

$$V_i' = \{T_i\boldsymbol{r}_j : q_1(\boldsymbol{x}_j) = \boldsymbol{c}_i^1\}. \tag{9}$$

We dubbed the generalized RQ model as transformed residual quantization (TRQ). While our generalization only requires transforms to be invertible, the property of non-increasing MSE error remain valid if transforms are orthogonal matrices, i.e., $(T_i)^T \times T_i = I$, (where I is the identity matrix) because applying orthogonal matrices does not change the MSE error. Yet, in this case, scaling inhomogeneity among residual clusters can not be accounted for by the transforms. Finally, denoting a reproduction cluster as $\widetilde{V}_i' = q_2(V_i')$, and Frobenius norm as $\|.\|_F$, the MSE for our TRQ model is given by

$$\text{MSE} = 1/N \sum_{i=1}^{k_1} \|T_i V_i - \widetilde{V}_i'\|_F, \tag{10}$$

Table 1: Complexity comparison with different VQ models, where m is the number of subspaces in PQ model, h is the number of levels in RQ and TRQ models, and l is the number of levels having transformations involved in our TRQ model; O_{dist} and O_{tran} are operation costs on distance computations and vector transformations respectively.

MODEL	# CELLS	ASSIGNMENT	MEMORY
VQ	k	kO_{dist}	kD
PQ	$(k^*)^m$	k^*O_{dist}	k^*D
RQ	$(k^*)^h$	hk^*O_{dist}	hk^*D
TRQ	$(k^*)^h$	$hk^*O_{dist} + lO_{tran}$	$hk^*D + lk^*D^2$

The complexity analysis of TRQ is presented in Table 1 along with the comparison with aforementioned VQ families. TRQ bring extra computational costs with transforms introduced in the model. In section 4, however, we will justify the advantage of TRQ in nearest neighbor search context with consideration of the substantial increase of search accuracy and reasonable control of using additional computational resources.

3.2. Iterative Residual Cluster Alignment

The remaining question is how to estimate T_i's that would be effective. An intuitive way is to align all clusters to a specific target set of points. The cluster V_i with the greatest population, for instance, is widely used in Procrustes shape analysis [10]. In this way, a one-to-one matching is required between the two sets, some extra computations then are performed due to the population difference. At the same time, you may have found that a one-to-one matching already exists between $\boldsymbol{r}_j$ and its reproduction $\widetilde{\boldsymbol{r}}_j = q_2(\boldsymbol{r}_j)$. Therefore, a more reasonable choice is a direct alignment from a cluster V_i to the reproduction set $\widetilde{V}_i$. A byproduct of this alignment is that it directly meets the ultimate goal to minimize the objective function in Equation 10.

Based on these observations, we propose an iterative alignment (IA) approach to finding T_i's. Where we iteratively update all the residual sub-codebooks in C^2 and the orthogonal matrices T_i's until MSE converges, or a specified number of iterations has reached. We formulate the two steps as follows,

Fix T_i's, update C^2: When T_i's are fixed, the procedure is the same as a normal iteration in a vanilla residual quantizer. There are two step involved:

1. updating the assignment I_j^2 for each $\boldsymbol{r}_j$, i.e. the indexes of centers $q_2(\boldsymbol{r}_j)$;
2. updating the codebook C^2 according to the new assignment $B = \{I_j^2 : j = 1, ..., N\}$.

Fix C^2, update T_i's: When the residual codebook C^2 is fixed, we need to find an optimal orthogonal matrix T_i for each residual cluster that

$$T_i = \arg\min_{\Omega} \|\Omega V_i - \widetilde{V}_i'\|_F, \tag{11}$$

where $\Omega^T\Omega = I$. The equation is solved by the Orthogonal Procrust analysis. The solution is straightforward. Given the covariance matrix $M = V_i\widetilde{V}_i'$, we use the singular value decomposition (SVD) $M = U\Sigma\ V^*$, to have the objective matrix $T_i = UV^*$.

3.3. Optimization to Inverted-Index Based ANN Search

The merits of RQ and PQ techniques for large-scale nearest neighbor search are two folds. For indexing, they produce very large number of landmark points that partition the space, and can be used with efficiency as hashing for fast indexing [27, 26, 8, 35, 36, 3]; On the other hand, they can produce a very compact representation of the vectors, and enable us to store huge amount of data in memory [19, 20, 11, 28]. Here we briefly introduce two state of the art ANN search systems, and then discuss optimization to the systems with our IA method. More information about parameter settings and experimental results are given in Section 4.

In the system introduced by Jegou et al. [19, 20], the first level is an indexer, which employs the k-means to construct an inverted file (IVF). In the second level, a product quantizer is used to compress the residuals to product codes. An approximate search then consists of a shortlist retrieval from the indexer IVF, and candidates re-ranking by the product codes. As database size increases, however, the inverted file becomes a bottleneck since the complexity for a k-means codebook learning is $O(Nk)$ and a vector assignment is $O(k)$. Where k is the codebook size or the number of cells in IVF. In order to address this issue, Babenko et al. [3] introduced the inverted multi-index to replace the IVF in above system. The new indexer employs a novel algorithm to jointly sort product cells in a PQ model. The complexities are reduced to $O(Nk^{\frac{1}{m}})$ for PQ codebook learning and $O(k^{\frac{1}{m}})$ for an vector assignment. It has been shown that both search speed and quality were improved. Ge et al. [11] recently applied their optimized PQ (OPQ) approach to the later system. The two levels are both optimized because they are essentially two PQ models.

Both of the systems are indeed based on hybrid models of residual product quantization. In the IVF system, the first level is a flat k-means, and the residual level is an ordinary product quantizer. At the same time, the inverted multi-index seems to be more complicated since two-level PQ models involved. However, we could simply see it as m_1 residual product quantizers independently produced from the m_1 subspaces in the multi-index structure. To optimize such residual product quantizers by our IA method, we only need to pay attention to the steps of the assignment and residual codebook C^2 updating, to make them follow the way of an ordinary PQ codebook learning.

4. Experiments

In this section, we examine the proposed TRQ model's performance on several large scale benchmark datasets. We consider both quantization distortion (i.e. mean squared error (MSE)) and ANN search. The first task is a direct measurement of the quantization error of different models, and the second task is a real world ANN search application of these quantization models.

Table 2: Statistics of datasets used.

DATA SET	# DIM.	# BASE	# TRAINING	# QUERY
MNIST	784	60K	60K	10K
SIFT1M	128	10^6	100K	10K
GIST1M	960	10^6	500K	1K
SIFT1B	128	10^9	10^8	10K

We have used four large scale datasets in our evaluation. The data types include images of hand-written digits from MNIST, local image descriptors SFIT [25], and global image descriptors GIST [29]. See Table 2 for statistics of the datasets. We will present in-depth experimental analysis on three relatively small datasets (MNIST [34], GIST1M, SIFT1M [19]), and report large-scale ANN results on the largest SIFT1B dataset [20].

In our experiments, we mainly consider two strong quantization baselines. The first method is product quantization (PQ) [19] which groups features into different groups, and performs quantization separately. The second method is optimized product quantization (OPQ) [28], which learns a rotation matrix for PQ to make feature groups more in-

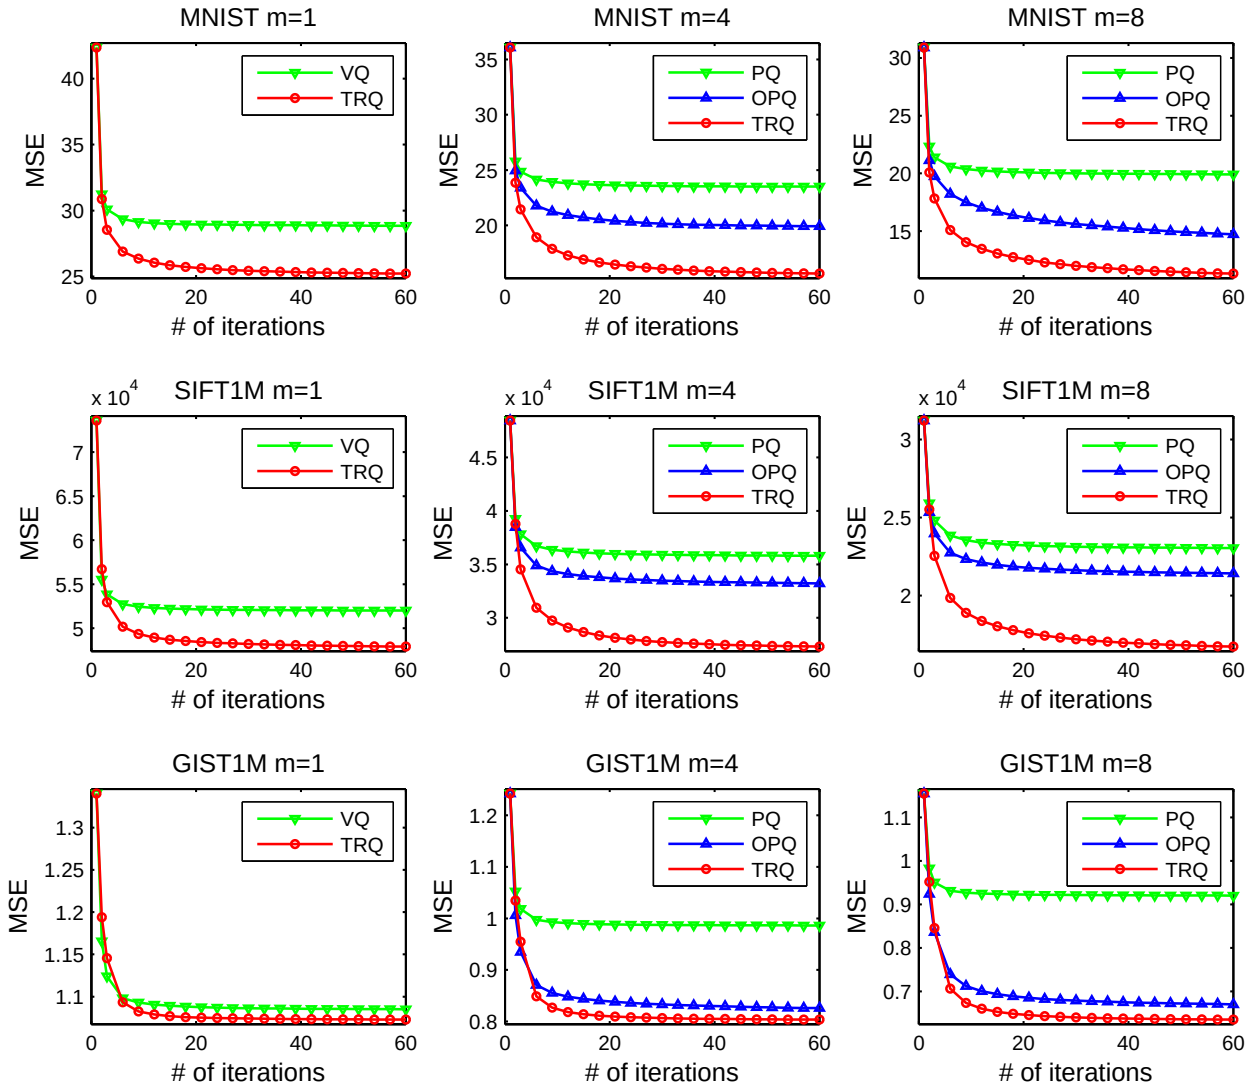

Figure 2: Comparative evaluation about distortion performance. The first level quantizers are fixed across all models so we only examine the behaviors on the residual quantizers. For SIFT1M and GIST1M, we use $k1 = 32$ and $k2 = 256$ that are numbers of clusters used in two levels. While for MNIST, $k1 = 10$ and $k2 = 64$. And $m-$ is the number of feature groups for residual quantizers. In addition, the codebook trainings are directly performed on the base sets.

dependent with each other, and to make quantization error smaller. This method has been shown to produce state of the art performance for ANN search.

4.1. Evaluation on Quantization Distortion

We first report quantization distortion for each model on different datasets, and we use mean squared error (MSE) as the distortion measurement. We only consider a 2 level quantization model in this section (no additional stacked layers for RQ and TRQ), which will help us better evaluate the quantization distortion. The parameter settings are given in Figure 2. For PQ, when the number of feature groups $m = 1$, it does not perform any grouping to the feature dimensions, and is reduced to an ordinary VQ (k-means clustering).

Figure 2 shows clearly that our proposed method TRQ has achieved a very significant gain over all other methods. Both TRQ and OPQ work better than the vanilla PQ model, which shows learning these optimal rotations has a clear advantage. Our method further improves OPQ by a significant margin. This is probably because our model does not make any independence assumptions about the feature dimensions or feature groups, while PQ and OPQ heavily

rely on this assumption, which is usually too strong for real-world data. It also shows that the gain of our methods is much more significant on MNIST and SIFT than that on GIST. This is probably because the feature dimensions of GIST are more independent, as has been shown and verified in [19]. In general, on all three datasets, our proposed TRQ method has achieved consistent and substantial gain over all other state of the art methods.

4.2. Evaluation on ANN Search

Our second experiment applies our method to optimize the ANN search system from [19]. The system is a typical two-level RQ model described in Section 3, which first applies k-means clustering and cluster the dataset into k clusters. All the residual vectors are then passed to a PQ to learn the compact codes. During a search, for a query point $\boldsymbol{y}$, we first compute its distance to all the first-level cluster centers, and probe $\boldsymbol{y}$ into the nearest w clusters, and then compute corresponding residuals $\boldsymbol{r}_j, j = 1, ..., w$ for $\boldsymbol{y}$. If $\boldsymbol{r}_j$ is from cluster i, then we need to rotate it as $\boldsymbol{r}'_j = T_i \boldsymbol{r}_j$ in our TRQ model. Finally, each $\boldsymbol{r}'_j$ is used to re-rank all candidates in its cluster through their PQ-codes. Depending on the dataset size, we use $w = 6$ for SIFT1M and GIST1M, and $w = 2$ for MNIST. We use Recall@R to measure the search quality, which represents the recall of the nearest neighbor within the top R candidates.

Figure 3 shows a similar trend as that of the previous section. Our TRQ method has achieved significantly better performance than the other two methods. For instance, on SIFT1M, we have achieved 31.51% on Recall@1, while OPQ achieved 25.77% and PQ 24.61%, respectively. Similar to the quantization distortion experiments, the gain of our method is much smaller on the GIST1M. This is probably because the feature dimensions of GIST are more independent (as verified by [19]). The true underlying reasons are to be further investigated.

4.3. Large-Scale Experiments with The Inverted Multi-Index

The last experiment shown in this paper is to integrate our TRQ model with the inverted multi-index [3] to improve its search accuracy on SIFT1B dataset. From Table 3, we can see that the search accuracy was substantially improved by our TRQ , which is up to 8% Recall@1 comparing to OPQ.

As the discussion in Section 3.3, the basic search steps are similar to the above IVF system [19], which include shortlist retrieval and candidates re-ranking by PQ-codes. We here refer interested readers to [3] and [11] for more details about the inverted multi-index and its optimized version called OMulti [11]. Here we are going to discuss more details about the search procedure when our TRQ integrated with OMulti, and more comparative results with other state of the art methods.

Table 3: ANN search comparisons between our model TRQ with OPQ [12] on SIFT1B with 16-byte codes (m=16) per vector. This table corresponds to Table 5 in [12]. T is the length of shortlist retrieved by the optimized multi-indexer OMulti [3, 12] for final re-ranking. The time is average search time per query that consists of shortlist retrieval and re-ranking.

METHOD	T	R@1	R@10	R@50	t(MS)
OPQ	10K	0.359	0.734	0.791	4.2
TRQ	10K	0.426	0.769	0.795	4.8
OPQ	30K	0.379	0.818	0.907	7.9
TRQ	30K	0.446	0.868	0.914	8.6
OPQ	100K	0.385	0.851	0.961	18.2
TRQ	100K	0.465	0.911	0.972	19.8

Parameters There are mainly four parameters involved, where m_1 and m_2 are used to represent the numbers of feature groups for the two different level PQ models, while k_1 and k_2 as the numbers of centers. We choose $m_1 = 2$ and $m_2 = 16$, with $k_1 = 16384$ and $k_2 = 256$, to be consistent with [3] and [11]. Since the multi-index has split the features to two groups (see $m_1 = 2$), the system can simply be seen as a concatenation of two independent residual product quantization models, and each of them has $m_1 = 1$ and $m_2 = 8$ while k_1 and k_2 are retained. In our case, they are optimized as two TRQ models.

From now on, we only need to present how one of the TRQ models works during a search. Note that, we are integrating TRQ with OMulti, so the whole raw space has been transformed by, say T^1, the transformation learnt in OMulti. A query $\boldsymbol{y}$ becomes $\boldsymbol{y}' = T^1 \boldsymbol{y}$ before it probes to the first-level clusters. We split y' equally to two segments, $\boldsymbol{y}' = \boldsymbol{y}'_1 | \boldsymbol{y}'_2$. Here we show only what happens to $\boldsymbol{y}'_1$; it then will be probed into the left subspace. If it is assigned to the sub-center $\boldsymbol{c}^2_{1i}$, in our setting, the residual for $\boldsymbol{y}'_1$ is $\boldsymbol{r}''_1 = T^2_{1i} \boldsymbol{r}'_1 = T^2_{1i}(\boldsymbol{y}'_1 - \boldsymbol{c}^2_{1i})$. Finally, the concatenation $\boldsymbol{r}''_1 | \boldsymbol{r}''_2$ are used for re-ranking PQ-codes on the shortlist that is provided by the indexer OMulti.

Computation-Wise Discussion A potential concern might be centered on the extra computational and storage costs required by the proposed TRQ model. TRQ applies transformations in the residual spaces and aims to improve the quantization quality and furthermore the search accuracy in the ANN problem. It slightly increases memory usage due to the local transformations. In the setting for Table 3, for instance, it requires about 2.5% extra memory space for all the transformation matrices relative to PQ or OPQ model. This is acceptable with the substantially increased search accuracy considered.

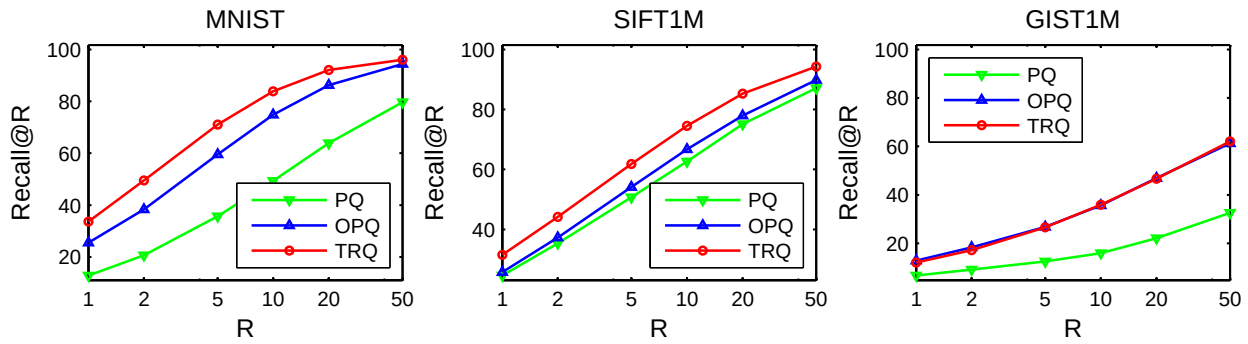

Figure 3: ANN search evaluation on different datasets. The parameters used are exactly the same as in Fig. 2, except that here $m = 8$ is fixed for all.

Even though our presented results have been restricted to comparisons with PQ and OPQ, we here give a short discussion about comparison between our method with LOPQ [21] considering the common feature of the locally learned transformations. LOPQ learn different PQ codebooks for each cluster while our TRQ share just one PQ codebook among all clusters. This results in LOPQ use much more extra memory than TRQ dose. In the same setting as Table 3 LOPQ produce 16384 PQ codebooks that takes 2 GB (about 10%) extra memory compared to ours and other methods. Taking advantage of the considerable memory usage, LOPQ have achieved about 2% higher recall than our TRQ on this SIFT1B dataset.

Computationally, the number of transformations or matrix multiplications in both TRQ and LOPQ for each query is determined by the number of visited cells from the first level index; it is usually less than 100 out of 16384 (in Table 2). The results in Table 3 show that the extra computational costs can be well justified given the significant improvements on search accuracy. We also use only 10 millions of the training points to speed up the codebooks learning in both of OPQ and TRQ. All the results have been produced using a single thread with an Intel(R) Xeon(R) CPU @ 3.40GHz with 25.6M cache size, and 128 GB RAM. All methods have been implemented in MATLAB and C (for codebook learning and database encoding) and C++ (for searching).

5. Conclusions and Future Directions

Product quantization and residual quantization are the well-known quantization models that reduce both computation and storage complexities. We have seen that different ways to optimize these models are becoming active and fruitful in the computer vision area. In this paper, we have presented a novel approach to reducing quantization error of residual quantization for fast nearest neighbor search. We have also presented two hybrid searching architectures that are based on product quantization and the proposed transformed residual quantization. Our experiments show that the proposed TRQ has achieved significantly smaller quantization error than previous methods, which clearly demonstrates the benefits of the proposed model. Also, in large-scale ANN experiments on 1 billion vectors, our method has achieved significantly better accuracy than previous methods.

Despite that we have restricted our attention to the ANN search problem in this paper, the approach can be extensively used for many other problems that require efficient and effective very large-scale clustering. For examples, kNN graph construction, and feature matching for image retrieval, and for object recognition to name a few. In addition, how to integrate our transformation-based approach with the global codebook learning approaches arises to be a very interesting question for residual quantization optimization. These become our future research directions.

References

[1] ANDONI, A. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. In *FOCS'06* (2006), IEEE Computer Society, pp. 459–468.

[2] BABENKO, A., AND LEMPITSKY, V. Additive quantization for extreme vector compression.

[3] BABENKO, A., AND LEMPITSKY, V. S. The inverted multi-index. In *CVPR* (2012), pp. 3069–3076.

[4] BARNES, C. F., RIZVI, S. A., AND NASRABADI, N. M. Advances in residual vector quantization: a review. *IEEE Transactions on Image Processing*, 2 (1996), 226–262.

[5] BOIMAN, O., SHECHTMAN, E., AND IRANI, M. In defense of nearest-neighbor based image classification. 1–8.

[6] CHAN, W.-Y., GUPTA, S., AND GERSHO, A. Enhanced multistage vector quantization by joint codebook design. *IEEE Transactions on Communications 40*, 11 (1992), 1693–1697.

[7] CHARIKAR, M. S. Similarity estimation techniques from rounding algorithms. In *In Proc. of 34th STOC* (2002), ACM, pp. 380–388.

[8] CHEN, Y., GUAN, T., AND WANG, C. Approximate nearest neighbor search by residual vector quantization. *Sensors 10*, 12 (2010), 11259–11273.

[9] DATAR, M., AND INDYK, P. Locality-sensitive hashing scheme based on p-stable distributions. In *SCG'04* (2004), ACM Press, pp. 253–262.

[10] DRYDEN, I., AND MARDIA, K. *Statistical shape analysis*. Wiley series in probability and statistics. Wiley, Chichester [u.a.], 1998.

[11] GE, T., HE, K., KE, Q., AND SUN, J. Optimized product quantization for approximate nearest neighbor search. *In CVPR* (2013).

[12] GE, T., HE, K., KE, Q., AND SUN, J. Optimized product quantization. *IEEE Trans. Pattern Anal. Mach. Intell. 36*, 4 (2014), 744–755.

[13] GERSHO, A., AND GRAY, R. M. *Vector quantization and signal compression*. Kluwer Academic Publishers, Norwell, MA, USA, 1991.

[14] GONG, Y., AND LAZEBNIK, S. Iterative quantization: A procrustean approach to learning binary codes. In *CVPR* (USA, 2011), pp. 817–824.

[15] HE, J., KUMAR, S., AND CHANG, S.-F. On the difficulty of nearest neighbor search. In *International Conference on Machine Learning (ICML)* (Edinburgh, Scotland, 2012).

[16] JAIN, A. K., MURTY, M. N., AND FLYNN, P. J. Data clustering: a review. *ACM Comput. Surv. 31*, 3 (Sept. 1999), 264–323.

[17] JÉGOU, H., DELHUMEAU, J., YUAN, J., GRAVIER, G., AND GROS, P. Babaz: A large scale audio search system for video copy detection. In *ICASSP* (2012), pp. 2369–2372.

[18] JEGOU, H., DOUZE, M., AND SCHMID, C. Hamming embedding and weak geometric consistency for large scale image search. In *ECCV* (Berlin, Heidelberg, 2008), ECCV '08, Springer-Verlag, pp. 304–317.

[19] JÉGOU, H., DOUZE, M., AND SCHMID, C. Product quantization for nearest neighbor search. *PAMI 33*, 1 (Jan 2011), 117–128.

[20] JÉGOU, H., TAVENARD, R., DOUZE, M., AND AMSALEG, L. Searching in one billion vectors: re-rank with source coding. In *ICASSP* (Prague Czech Republic, May 2011). QUAERO.

[21] KALANTIDIS, Y., AND AVRITHIS, Y. Locally optimized product quantization for approximate nearest neighbor search. In *CVPR* (Columbus, Ohio, June 2014), IEEE.

[22] KULIS, B., AND GRAUMAN, K. Kernelized locality-sensitive hashing for scalable image search. In *ICCV* (2009).

[23] LIU, T., MOORE, A. W., GRAY, A. G., AND YANG, K. An investigation of practical approximate nearest neighbor algorithms. In *NIPS* (2004).

[24] LIU, W., WANG, J., JI, R., JIANG, Y.-G., AND CHANG, S.-F. Supervised hashing with kernels. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2012).

[25] LOWE, D. G. Distinctive image features from scale-invariant keypoints. *IJCV 60* (November 2004), 91–110.

[26] MUJA, M., AND LOWE, D. G. Fast approximate nearest neighbors with automatic algorithm configuration. In *ICCVTA, VISSAPP'09* (2009), pp. 331–340.

[27] NISTER, D., AND STEWENIUS, H. Scalable recognition with a vocabulary tree. In *CVPR* (2006), vol. 2, pp. 2161–2168.

[28] NOROUZI, M., AND FLEET, D. Cartesian k-means. *In CVPR* (2013).

[29] OLIVA, A., AND TORRALBA, A. B. Modeling the shape of the scene: A holistic representation of the spatial envelope. *IJCV 42* (2001), 145–175.

[30] PHILBIN, J., CHUM, O., ISARD, M., SIVIC, J., AND ZISSERMAN, A. Object retrieval with large vocabularies and fast spatial matching. In *CVPR* (2007).

[31] TORRALBA, A., FERGUS, R., AND WEISS, Y. Small codes and large image databases for recognition. In *CVPR* (2008).

[32] WANG, J., KUMAR, S., AND CHANG, S.-F. Semi-supervised hashing for scalable image retrieval. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (San Francisco, USA, June 2010).

[33] WEISS, Y., TORRALBA, A., AND FERGUS, R. Spectral hashing. In *NIPS* (2008), pp. 1753–1760.

[34] YANN LECUN, CORINNA CORTES, C. J. B. MNIST. `http://yann.lecun.com/exdb/mnist/`, 1998. [Online; accessed 20-May-2014].

[35] YUAN, J., AND LIU, X. A novel index structure for large scale image descriptor search. In *ICIP* (2012), pp. 1937–1940.

[36] YUAN, J., AND LIU, X. Transform residual k-means trees for scalable clustering. In *ICDM Workshops* (2013), pp. 489–496.

[37] ZHANG, T., DU, C., AND WANG, J. Composite quantization for approximate nearest neighbor search. In *ICML* (2014), T. Jebara and E. P. Xing, Eds., JMLR Workshop and Conference Proceedings, pp. 838–846.